

8 Nonidentification

Identification like scars or vapor trails arising and disappearing [#1585]

What exactly does it mean for consciousness or awareness to be “pure”? There are all kinds of places we can start if we want to answer this question. People often start with ancient religious belief systems, or with grand philosophical theories of what the purported “essence” of consciousness is. But even at this early stage, it is already becoming clear that if we start from the real-life phenomenology of meditative experience itself, there is more than one answer to our question. The essentialism of looking for or expecting a single answer may itself be the problem, and it may be that to make progress, we need to let go of the hope for the one big answer. Multiple answers do not necessarily exclude each other; they may even complement each other in forming a bigger picture. Let me therefore begin by slowly flagging a series of possible answers to the question of what the adjective “pure” in expressions like “pure consciousness” and “pure awareness” could actually mean. Many of our reports talk about contentlessness. Traditional sources do too. The first, classical and quite radical, reply is: “Consciousness is pure whenever there is no other experiential content whatsoever, when the quality of *consciousness* itself is the only kind of phenomenal character that can later be remembered and reported.” This is the kind of pure consciousness that we find in full-absorption episodes. In the introduction of this book, we saw that there are two main categories of such episodes: in deep states of meditation arising out of waking consciousness, and in dreamless sleep (see chapter 20). Here, a scientist might say that in these cases, minimal phenomenal experience (MPE)—the simplest, minimal form of experience—occurs as a singular, stand-alone feature. This is to say that in these episodes, there is no content, just naked awareness itself.

The phenomenological question of what the “purity” of pure consciousness actually consists in has a second, equally canonical reply. This one associates the “purity” with

phenomenal qualities like “peace” and “deep, unbounded silence,” the absence of all mental conflict, noise, and perturbation. We first encountered support for this view in chapters 2 and 3. Our phenomenological data confirm what has been known for millennia: Consciousness can exist entirely without thought. On a simplified reading, pure conscious experience here means the absence of all discursive thought, including memory, planning, daydreaming, and mind-wandering. For example, one well-known formula is that pure awareness is “transcendental consciousness” in the sense of a state that has “transcended thought,” where we simply abide in a crystal-clear and thoughtless state—the state of knowing restful alertness itself.

The third reading of “purity” is “clarity” or “clear wakefulness,” the experience of a clear and unobstructed space of knowing. This answer can also be found in our data, as well as in some traditional frameworks. We have already looked at this specific experiential quality in the previous two chapters, where I introduced the concept of “epistemic openness” (to help explain wakefulness) and the idea of having an inner “model of an unobstructed epistemic space itself” (for clarity). This version of “purity” is the nonconceptual conscious experience of one’s own epistemic capacity *as such*, a space of *potential* states of knowledge, without any priors or predictions; an inner space that is wide open and lucid at the same time. Here, consciousness is pure in the sense of being “capacious,” being a lucid, open space in which experience can occur. The space is not clouded by thoughts.

In chapter 7, we even encountered a radically naturalist reading: Hypothetically, pure awareness could be a way of thoughtlessly experiencing the activity of a specific, as yet unknown part of the *neural* body only.¹ This fourth version would be “pure” in not being mediated by any sensory system (either directed at the interior of the body or aimed outward into the environment) and in emerging without any form of inner or outer action. From a neurocomputational perspective, we could say that pure awareness, in this sense, is special because it arises neither from perceptual nor from active inference. It would be the most abstract experience of embodiment that human beings are capable of (see chapter 24 for more). As shorthands for these options for making sense of pure awareness, we could even begin to speak of P1 (no content), P2 (no thought), P3 (clarity), P4 (abstract embodiment—here viewed only as a computational property, but see chapter 24), and so on. In this chapter and the next, we will encounter two further phenomenological interpretations of the term “pure consciousness.” I promise that I will give a full list of these in chapter 34 of this book, when summarizing the main results of our investigations. But the beauty of lived contemplative experience is that it does not adhere to any rigid conceptual schema.

Phenomenologically, our first three interpretations of “purity” can clearly coexist in one experience. For example, during meditation practice, it sometimes happens that one enters a deep state of full absorption that is characterized by nothing other than a crystal-clear experience of epistemic openness and unbounded, wakeful silence (P3). But in real life, all these aspects tend to get mixed up. P3 and P2 may give way to P1, leading into a process in which the practitioner oscillates back and forth between thoughtless clarity and full absorption multiple times, in a way that is extremely hard to describe in words. However, sooner or later, every meditator will realize that this is not the only way of experiencing pure consciousness. And this is where things get interesting. What happens when you sit in a deep and stable state of restful alertness and very slowly open your eyes, now seeing the world with a clear and thoughtless mind? What exactly happens when, while in a state of reflexively aware mindfulness, you hear sounds emerging out of silence? Phenomenologically, are the sounds perhaps sometimes *made* out of silence? And what happens if, in that crystal-clear silence, one single thought arises, unfolds its content, and quietly dissolves, as in Padmasambhāva’s classic image of a thief entering an empty house, looking around, and disappearing?

Padmasambhāva came to Tibet around 767 and is widely venerated as a second Buddha, even in Nepal, Bhutan, and the Himalayan states of India. He said that just as an empty house is at no risk from a thief because there is nothing that the thief could grasp and hold on to, thoughts recognized as empty cannot in any way harm an empty mind. Just like the thief leaving the house, they will “self-liberate” and quietly dissolve. We were never taught to look there—but is it possible that thoughts are really *made of* that self-knowing mindfulness itself? Padmasambhāva called the state resulting from this non-intellectual insight a state that had “crossed the dangerous defile of moving thoughts.”

The main point is that the openness of pure awareness can become interestingly intermingled with various modes of conscious experience. One of our participants beautifully described the coemergence of a stable state of pure awareness and conscious content by creating the following two metaphors: “A gentle rain on a still pond, or a crystal ball with water being poured over it” (#1293). And of course, four modes are not all there is. Another—the fifth, and equally valid—reading of what the “purity” of MPE amounts to is this: The distinct and timeless quality of awareness itself can be clearly felt, the phenomenal character of epistemic openness exists, but there is additional content that in principle could be reported (e.g., spontaneously arising thoughts, visual perception, bird song, bodily sensations, etc.), and yet the process of *identification* is lacking. The thief doesn’t try to grab anything. There is no mental agent attempting to control the process. The meditator has taken a break. Now, conscious experience

may even be pure in the sense of having become “nondual,” lacking subject/object structure (chapters 26 and 27). And what is seen when the eyes open, then, may be neither real nor unreal (more on the phenomenology of “virtuality” in chapter 28). Here, the experiential quality of awareness is pure in the sense that it remains autonomous, uncontracted, and fully present, but *detached* from any spontaneously arising content. To put it in traditional Buddhist terms, grasping and clinging are absent. To use more modern terminology, from current philosophy of cognitive science: Internal representations of ongoing cognitive or perceptual processes are not automatically integrated into a transparent self-model anymore.² Ownership and the sense of control are absent. Therefore, another phenomenological reading of “pure consciousness,” can be summarized as “no identification.” Let us look at some experiential reports that illustrate this version:

182 I experienced this awareness as a distinct state of consciousness that came on its own and after a while (many minutes, not many hours) passed away again. I perceived myself as completely detached from my thoughts and sensations. These still existed and I still perceived them, but was no longer identified with them. Rather I was identified with an awareness that perceived these thoughts and sensations without feeling anything myself. In this respect I would describe it as pure.

1703 [. . .] It starts when I look for “who” is doing the seeing. Or “who” is doing the feeling? [. . .] There is certainly physical sensation happening. That is very clear. But who is feeling the feeling? Or is there just feeling? Or that same line of thought, but with hearing or thinking. As soon as “I” look for the thinker of thoughts, it hits me: There is simply thinking. But no one or thing is doing it. It’s just happening all on its own. And it’s like there is nothing behind the experience. There is just sort of floating. Just like a big smear of sensations all suspended somewhere. And then as quickly as it comes, it goes . . . in the sense that, thoughts of how “I” just had a cool experience and of ways to get it back. I start clinging. I start wanting it to come back and stay for longer. . . . When I’m having the experience, certain bits of language seem completely nonsensical. For example: “I hear a sound.” It seems that “There is hearing going on” feels so much more of a correct description. The same goes for, “I am thinking a thought” → there is thinking going on.

One particularly relevant phenomenal quality for the theory and practice of de-identification is the “sense of control,” the subjective experience of being the doer of actions and the thinker of thoughts:

1830 [. . .] Acceptance of not being “the thinker” was far more accepted and clearer than usual. Awareness had space for everything, while not “being” any single thing. Breath proceeding normally and being known, but in larger space than usual, perhaps unbounded space. Peaceful yet not emotional. Neither filled with agency nor completely agencyless. Some relief from not having to construct/feel in control.

Nonidentification also applies to soundness and the affective dimension of MPE experiences:

246 It was pure joy without attachment.

2524 To everything that is thought and imagined, one feels a connection (not in the sense of dependency or the like, but rather as a being-interwoven, as love) and yet clarity (in the sense of a kind of distance, not attracting, not identifying).

The phenomenology of spatiality (see chapter 23) in this version of purity is inversely related to mental grasping:

1819 [. . .] The awareness is very broad, beyond the body or spatial boundaries. Yet bodies or space appear in it—completely relaxed. Depending on the situation, the experience of spatiality can spread out strongly and increase in luminosity (not visual) and in power / luminosity / depth / vastness / clarity / knowledge, depending on how deep the connection is, depending on how little distraction / grasping there is.

As many meditators know, the process of carefully but ever more gently and effortlessly observing the contents of self-consciousness can gradually lead to an experience of de-identification. The phenomenological elements of control, mental agency, and selfhood begin to disappear, while an additional quality of clarity emerges. P3 and the quality of nonidentification are related. For example, there is a distinct and specific phenomenology of seeing thoughts arise within pure awareness itself and later dissolve back into it. This is the experience of *cognitive* de-identification, the conscious experience of naturally and gradually disengaging from the contents of thought. Now these contents are still conscious, but they are no longer part of the egoic self-model. The same can happen for emotions and bodily sensations.

Interestingly, nonidentification is reversible. Nonidentification can be reversed through the automaticity of “clinging,” of gradually returning to the phenomenal character of agency (as in the last example presented here). It can also, however, be reversed seemingly at will, and for reasons other than anxiety or unease, in a more relaxed,

exploratory fashion that nevertheless leads back to identification. To judge from these reports of real-life experience, something in us may also “deliberately” cause a reidentification with the content from which awareness was previously detached:

207 I was sad. On the way home from the bus I first tried to observe the sadness while walking and then to detect the consciousness that feels/observes the sadness. Suddenly I was in a state of absolute silence and peace that I had never experienced before. The sadness was still there, but it was no longer part of me and somehow had nothing sad about it anymore. I can't describe it, it was as if it was outside of me. Then I started to marvel at the situation and then my mind came back with the thought that I would always meditate so much and that I could just be sad for today. And in that moment it got me and I opted for sadness and the silence was gone.

Our phenomenological data also show that nonidentification can refer to single mental states, but also to the person as a whole:

2953 [. . .] Yet, I felt just as being myself, but not being anything in particular including being a human or individual person. In fact, I really didn't even know I was in such a state, yet being in it, until the mind suddenly produced a few thoughts, which I observed as if they were a phenomena outside myself. I could hear the quality of them as if being part of a person that I recognized, but no longer identified with.

Sometimes experiences like these are later described as coming into contact with the “true self” (chapter 29). Phenomenologically, de-identification can also lead to de-immersion. One classic metaphor for this process, found in many places in the popular literature on meditation, is the image of being fully immersed in a movie and then suddenly realizing that you are *not* the hero, disengaging from the perspective from which the story seems to be told. Here is one example from our study:

26 The first experience to which your description of “consciousness of consciousness” would apply happened when I was still very young. [. . .] It was in that moment that I felt for the first time that I am not my thoughts, but the invisible thinker. I felt like a lonely spectator in an empty cinema who has suddenly realized that he's not the lead actor in the film he's watching.

In this report, there is still a thinker; there is a metaphorical “lonely spectator,” although he is “invisible” and no longer fully immersed in his thoughts. If we were to feign ignorance and take the words at face value, then the description here would seem to refer to an example of dual meta-awareness, because the speaker still uses the

first-person pronoun “I” and there is a subject/object structure (a “lonely spectator” and the “thoughts”). However, it might well be that the actual episode of “awareness of awareness” was entirely nondual, but for lack of better words, it is later referred to with the help of a dualist metaphor (see chapters 26 and 27). We can also imagine another scenario. Here is one possible variation on our classic metaphor:

Imagine the cinema is dark and completely empty. All the seats are empty. Running on the screen is a movie of someone watching a movie, a person who is deeply immersed in the plot of this movie-in-the-movie most of the time, but who sometimes has glimpses and vague intuitions that she might actually not be the protagonist, but might instead be a passive viewer in some sort of cinema. Or that she might even be the whole cinema itself. Not seated in the middle, but somehow everywhere. From time to time, the viewer in the movie even has fleeting and slightly complacent philosophical fantasies that she might be some sort of empty space in which the movie becomes aware of itself, or that she might be the beam of light coming from the projector. Suspended in the dark, as it were, this movie containing the movie-in-the-movie is the only thing that can be seen. But nobody is watching. Now, very slowly, the light is being turned on, in slow motion as it were. The cinema as a whole emerges out of the darkness, slowly fading in, empty seats and all. Now the lights are on. The movie is still running. Nobody sees it.

The Contraction Principle

Of course, ultimately the little one is not real. He is a picture. He is not aware.

—Douglas E. Harding (1909–2007), *Face to No-Face*

If you have a background in philosophy or consciousness studies, you may remember those boring old discussions about zombies—round and round the question of whether there could be a creature behaving exactly like you or me but not conscious. As a way to start thinking about the contraction principle, let’s take the idea one step further, moving beyond mere intuition-mongering toward something more meaningful. If zombies are possible, it is possible that you are now unconscious—but that is impossible for *you* to discover this fact intellectually. A zombie is a functional isomorph of a possibly existing conscious being, which means that all its outer and inner behaviors and functions are indistinguishable from such a being. Not only does it walk and talk in the same way as its conscious twin, but it also thinks the same thoughts (e.g., when it thinks that it is actually conscious right now). Again, if zombies are possible, then you might be one

right now and you would have no chance of ever discovering that many of the beliefs that you have about yourself are actually false—for example, the belief that you are not a zombie and are quite obviously conscious in this very moment. Why? Because if you did, you would not be a functional isomorph of your conscious twin anymore. Your brain couldn't have the same causal structure—because it has just generated a thought that would have been impossible with the conscious twin's causal structure. All that any of this really tells us is that absurdities easily accumulate around conceivability arguments and intuition pumps in the philosophy of consciousness.

Given this background, let me now draw your attention to something much more interesting. The phenomenology of some spiritual experiences gives a new and much deeper meaning to thought experiments of this kind, as well as to other ideas, like the Douglas Harding quote presented at the start of this section, or the contraction principle (to be explained next), or the suggestion that no character in any narrative can ever be a truly self-aware part of the narrative (even if it tried to control the script and “wake up” to its own fictitiousness; see chapter 17). What we took to be ourselves—the content of our conscious self-model—could actually be a zombie. Our ordinary conscious self-model could ultimately be a mere image, a transparent image of an ego that falsely believes it is actually conscious. As Douglas Harding says, the “little one” may not be aware. What spiritual traditions sometimes call “insight” or “liberation,” then, could be the moment when the whole cosmos (not the egoic self) realizes that it is the entity that is conscious. This realization is not something intellectual, and not something that *you* (or the zombie) could ever have. The phenomenological result of this shift into the zero-person perspective would be the emergence of an open self-aware field, a globalized and nonegoic form of self-awareness permeating everything else. And this new model of reality would not merely be a new “belief”; it would constitute an epistemic gain of an entirely different kind. There are forms of conscious knowing that have nothing to do with some fictitious “self” being related to words, concepts, or propositions. As a matter of fact, it may turn out that beliefs are something that only zombies have.

Carefully investigating the phenomenology of pure awareness can thus lead to interesting discoveries. The phenomenology includes not only the way in which pure awareness emerges, but also the way it breaks down. Perhaps you have already discovered how, phenomenologically, it is not *we* who get identified or fused with some content (e.g., an arising thought)? In my own experience, I find that it is *awareness* that gets identified. “We”—as the agentive sense of self—come into existence only *after* the fact, after awareness has already contracted into a specific content (like the mental image of a future goal state to be reached, a subtle feeling of mental effort, the ensuing sense

of control, etc.)—after, as it were, open awareness has already been lost to itself. The sense of self is the *result* of contraction, not its origin. The breakdown is not your fault. You are not the cause, but the effect. You are the way in which the brain explains away something that was slightly unexpected. That “something” is the surprise of suddenly being able to control a part of the world—your body, your attention, your thoughts.

The “contraction principle” is easy to understand intellectually but very hard to investigate within phenomenal experience itself. From a scientific perspective, “being conscious” or “appearing” is a property of a complex, internal model in the brain: What really *is* conscious is simply a certain part of an organism’s model of the world, a specific processing layer in its internal model of reality—which typically also includes the organism itself plus other agents in the world. We have not yet understood what this property is, but it is straightforward and empirically plausible to say that “phenomenality” or “appearance” is a property of some sustained, functionally integrated state in our head. Consciousness is something subpersonal. At the very least, this is a rational and evidence-based theory, one possible view from the outside.

But seen from the organism’s inner perspective, things are very different. The model is transparent; therefore, the representational medium is invisible. Yet this model is all the organism has. This means that what it experiences is experienced not as a world-*model* or some sort of inner image, but simply as the world itself. Consciousness is the appearance of a world; there is full immersion, plus the phenomenology of direct realism. In addition, because the internal model is not being experienced *as* a model, it now is the *organism* that seems to be conscious, not the world as a whole or some complex image in the organism’s head. Phenomenologically, the property of “being aware” has been contracted. Now, consciousness is something personal. Now *you* are conscious.

From an outside perspective, however, it is as if the embodied brain—in searching for a viable strategy of portraying reality—found it helpful to commit what in philosophy is known as the “fallacy of composition.” This is an informal fallacy that arises in natural language when one falsely infers that something is true of the whole from the fact that it is true of some part of the whole. Two examples would be “Atoms are not alive, all biological creatures are made out of atoms, so ultimately none of them is really alive!” and “All voters have rational preferences, therefore any collective choice induced by majority rule will also be rational!” Of course, the brain evolved not to perform logic but to create an efficient predictive model of reality that helps the body to survive and copy its genes to the next generation. According to the virtual reality (VR) in your head, there is now a conscious self—a whole, embodied person—and this is what forms the origin of a first-person perspective.

Alongside the conscious self, there will often be other self-aware agents that are also portrayed as conscious and able to control their own attention, their thoughts, and their bodily movements. They are not empty persons,³ but egos. Phenomenologically, they are not dream characters or conspecifics with some complex VR running in their brains, but full-blown agents—and they constantly project this property back into us, as we begin to mirror each other. To cooperate successfully, we need to navigate a complex social world. From the organism's perspective, there is a social frame of reference, and the property of "being conscious" may now be instantiated at *multiple* locations in this environment. Even in the dream state, neurotypical human beings will always experience other dream characters as also "being conscious." The illusion—which often persists even during lucid dreams—that dream characters have a conscious life of their own illustrates a robust coding principle of the human brain: For the organism, "phenomenality" is a local, personal-level property either of itself or of another creature. It is not something either omnipresent or local. It is a global feature of the animal as a whole, and it can occur in multiple agents at the same time.

Philosophers of consciousness distinguish between system consciousness (whenever the predicate "conscious" is applied to an animal, a person, or a machine as a whole) and state consciousness (all cases where "conscious" is used to refer to a state of an animal, a person, or a machine).⁴ In ordinary states, the brain depicts consciousness as system consciousness, not state consciousness. Internally, a state property is *misrepresented* as a system property. Let us call this the "contraction principle":

(CP) "Phenomenality" is a subpersonal property, a property of certain functionally integrated brain states. The brains of neurotypical human beings *misrepresent* this objectively given property of phenomenality by contracting it into a transparent conscious self-model, which then forms the origin of a first-person perspective.⁵

And here is where research into the phenomenology of pure awareness gets really interesting. This research is directly relevant to constructing a minimal model of consciousness, and it will also be crucial in the formulation of a first standard model. Why? First, the phenomenological degree of "aperture" of the field of awareness is something that can be measured by meditation researchers.⁶ More important, there are well-documented counterexamples to CP, such as the nondual states of awareness that sometimes spontaneously occur in advanced practitioners of meditation. We will carefully look at such states later in this book, mostly in chapters 25, 26, and 27. In such nondual states, we find conscious experience without contraction, phenomenality without an egoic self-model, and awareness without the perspective created by top-down superimposition of an abstract subject/object prior. The organism is liberated

from full immersion, and it can transcend the phenomenology of naive realism. In such states, it is as if the perspectives of science and contemplative practice begin to meet, as if they might become congruent—but in an unexpected way, perhaps in a way that some may even find uncomfortable. What the contraction principle shows is that nondual states of MPE may actually be the best entry point for the formulation of a first standard model of consciousness. The minimal model approach dissolves the problem of subjectivity.

MPE research is essential because there is a methodological primacy to subpersonal state consciousness: It can occur without system consciousness, but not vice versa. This is shown by our phenomenological material, as we will see in the course of this book. Wakefulness, reflexivity, and the phenomenal signature of knowing can all exist in nonegoic variants, but egoic self-awareness is always accompanied by the experience of wakefulness, self-directedness, and a knowing self. What if you even *were* self-knowing wakefulness that had mistakenly identified itself with the epistemic agent model of the biological organism in which it occurred (chapter 24)?

Possibly your mother and your father played a role in this process. Egoic self-awareness is what makes social cognition possible in the first place. Please note that what I have termed “contraction” is a spatial metaphor for a form of misrepresentation that, in the wider context of biological and cultural evolution, has apparently been a successful form of self-deception: A property of a part was portrayed as a property of the whole; a subpersonal state was depicted as a personal-level property. In this wider context, there is nothing wrong with that. Yes, it has generated a lot of psychological suffering, but at the same time, one can view it as a major triumph of natural evolution. Imagine that you are a conscious animal living in an increasingly complex social environment populated by other such animals. They could be potential allies, they could be prey, and they could also be dangerous. It is important for you to know whether other such animals are currently responsive, whether they are asleep or awake, and whether they can attend to you or recognize you. It matters whether they are epistemically open to the world (chapter 4) and whether they know that they are (chapter 5).

This is where your parents come back in. Maybe as an infant, you were much more like a loosely structured field of open awareness, but Mom and Dad—who were already firmly in the grip of their own illusion of conscious agency—continuously projected an epistemic agent model onto you by stubbornly treating you as a single conscious entity, a knowing self. Later, through “education,” and via punishment and reward, they even held something that they called “you” responsible for outer behavior, stubbornly and at times cruelly assuming the existence of stable egoic self-awareness, of accountability and free will. Why did your parents do this? One may speculate that

“awareness contraction” was a necessary precondition for social cognition and cooperation in large groups of human beings. High-level social cognition made it useful for biological organisms to mutually represent each other as self-aware agents. You solve the problem of detecting and assessing degrees of epistemic openness by classifying other organisms that you perceive from the outside as either conscious or unconscious, or alternatively by attributing the graded property of “being aware” to them, thereby experiencing them as currently being more or less alert, more or less wakeful, more or less distracted, and so on. And then you turn the classifying impulse (whether binary or graded) back on yourself.

Awareness contraction is how you learn to perceive yourself as an individual, self-aware agent, an active entity that is currently conscious. And this may be one of the many ways in which pure, nondual awareness became contracted into a self-model: Given the single-embodiment constraint (more on this in chapter 27), the subject/subject distinction you *had* to use when interacting with members of your own family or tribe became the blueprint for the phenomenology of subject/object duality. Something that in reality is wide open, entirely passive, clear, and silent turned into the property of an active self via (1), most importantly, the single-embodiment constraint; and (2) the need for social transactions that depend on a subject/object split. Reflecting on these two reliable drivers of contraction, it is interesting to note how many of humankind’s most earnest practitioners of meditation have chosen to live in solitude as hermits, or in communal seclusion as nuns and monks, and how, to this day, silent meditation retreats take place in an environment in which all social interactions are deliberately reduced to a minimum.

More immediately, the contraction principle is important if we want to better understand the phenomenology of identification that is the focus of this chapter. Why are some features of awareness, even the phenomenal quality of awareness itself, almost automatically attributed to a fictitious self? Why are they integrated into a transparent self-model? From a scientific perspective, neither the self-model itself nor the person as a whole is conscious—the subpersonal world-model in the person’s head is what is truly conscious. This leads to a testable phenomenological prediction: If a meditator were to become gradually aware of this hypothetical fact, then this process should have two components. First, awareness that the world-model is the conscious entity should create a global phenomenology of nondual wakefulness, which we might (drawing on the title of David Hinton’s 2019 book) call an “awakening cosmos,” an awareness of *everything* becoming nonegoically aware and epistemically open. Second, this awareness should lead to a phenomenology of virtuality, to a gradual suspension of naive realism (more on this in chapter 28). This experiment would be fascinating to perform.

In the meantime, however, we might still find ourselves asking: But why was the identification with an agent, with a conscious self, so very successful?

For a biological organism and its brain, it is of the utmost importance to reliably distinguish between those events that are self-caused and those that are “external,” in the sense of currently being beyond our control. There now is a strong consensus across many disciplines in the consciousness community that neither the brain nor consciousness evolved to meditate, to think logically, or even to deliver an accurate inner image of the world as it “really” is. The primary goal is to continue staying alive, to maintain physiological integrity in the face of danger and opportunity, even if this should require clever forms of self-deception. In the words of Anil Seth:

Evolution’s reason for providing organisms with brains is not so they can write poetry, do crossword puzzles, or pursue neuroscience. Evolutionarily speaking, brains are not “for” rational thinking, linguistic communication, or even for perceiving the world. The most fundamental reason any organism has a brain is to *help it stay alive*, through making sure that its physiological essential variables remain within the tight ranges compatible with its continued survival.⁷

Conscious experience is about control. It is about predictability and reducing uncertainty. For many organisms, staying alive presumably involves successfully distinguishing between what they can control and what is beyond their reach. This is a fundamental categorization, leading much later to conceptual distinctions like “self” versus “nonself.” The first distinction (controlled/uncontrolled) then turns into a “prior”: a deeply ingrained and very successful top-down prediction about what the world is like and how it will be experienced. A self exists, a single “distinguished entity” that is the locus of control and constitutes the experiential unit of identification (see chapters 24 and 29 for more).⁸ The rare phenomenology of nonidentification is what suspends this distinction. It is not that the meditator falls into an explicit phenomenology of alienation, of being externally controlled (as the schizophrenic does). It is rather that the distinction itself gently and quietly disappears—but without the accompanying thought, “The distinction between self and nonself has disappeared!” All that happens is that one self-fulfilling prophecy turns silent. In terms of a computational model, we could say that the anticipation, the top-down prediction that every perceptible event must be either self-caused or external, is suspended. Within the experience, the distinction becomes meaningless, as in report #1830: “[. . .] Neither filled with agency nor completely agencyless. [. . .].” It is as though an abstract hyperprior has been episodically suspended. What results is a liberating form of indeterminacy, a phenomenology of neither-nor-ness.

But there is more. If it is true that the sense of self is the *result* of contraction, not its origin, then all of this applies to the “meditating self” too. If our unfolding story is on the right track, then there should actually be very subtle illusions of control. If you are a practitioner yourself, can you observe what exactly happens in the very moment when you notice that your mind has wandered, that you have “lost it again”? The meditating self is the new phenomenological entity that is disappointed with itself, the noble one that sincerely wanted to fully identify with pure awareness, but somehow “lost control over the process.” That wasn’t its fault. There is no reason to be disappointed.

Unfortunately, the phenomenology of “regaining” control and mindfulness, of “waking up again,” could often be an example of exactly the same mechanism. No reason for the meditating self to be proud, to pat itself on some invisible shoulder with a feeling of success and relief. The subtle psychological drama, the self-condemnation, the relief, and even the practitioner’s sense of achievement, stamina, and sustained discipline could always be just the result of yet another contraction. Occasionally, the whole dynamic can be seen through (though I rarely manage to do that myself).

Given our new conceptual instrument of “contraction” as a special form of misrepresentation, specifically a biologically successful computational principle, we can now describe the phenomenology of meditation more precisely and begin to distinguish different ways in which the quality of awareness per se may manifest. Let us briefly review the preceding five chapters through the lens of the contraction principle. We have seen that there are contracted and uncontracted forms of silence because the phenomenology of silence can be experienced as being “in the mind,” but sometimes also as all-pervading. Wakefulness can be a property of the meditating self or of the world as a whole (“the awakening cosmos”). Similarly, the phenomenology of clarity can be experienced as a local form of lucidity in one’s “own” mind, but it can also expand and sometimes even turn into a feature characterizing the space of conscious experience as a whole. The density of a “thick silence” can be located in the head or it can be everywhere. Soundness and internal harmony are often experienced as related to an embodied self, but sometimes also as qualities of the overall situation, pertaining to the wholeness of the moment. Let us keep this new conceptual tool of contraction in mind as we penetrate deeper into the landscape of seeing *what is*.

This is a section of [doi:10.7551/mitpress/15196.001.0001](https://doi.org/10.7551/mitpress/15196.001.0001)

The Elephant and the Blind

The Experience of Pure Consciousness: Philosophy,
Science, and 500+ Experiential Reports

By: Thomas Metzinger

Citation:

*The Elephant and the Blind: The Experience of Pure Consciousness: Philosophy,
Science, and 500+ Experiential Reports*

By: Thomas Metzinger


DOI: 10.7551/mitpress/15196.001.0001

ISBN (electronic): 9780262377287

Publisher: The MIT Press

Published: 2024

The open access edition of this book was made possible by generous funding
and support from MIT Press Direct to Open



The MIT Press

© 2024 Massachusetts Institute of Technology

This work is subject to a Creative Commons CC-BY-ND-NC license.

This license applies only to the work in full and not to any components included with permission. Subject to such license, all rights are reserved. No part of this book may be used to train artificial intelligence systems without permission in writing from the MIT Press.



The MIT Press would like to thank the anonymous peer reviewers who provided comments on drafts of this book. The generous work of academic experts is essential for establishing the authority and quality of our publications. We acknowledge with gratitude the contributions of these otherwise uncredited readers.

This book was set in Stone Serif and Stone Sans by Westchester Publishing Services.

Library of Congress Cataloging-in-Publication Data

Names: Metzinger, Thomas, 1958– author.

Title: The elephant and the blind : the experience of pure consciousness: philosophy, science, and 500+ experiential reports / Thomas Metzinger.

Description: Cambridge, Massachusetts : The MIT Press, [2024] | Includes bibliographical references and index.

Identifiers: LCCN 2023012135 (print) | LCCN 2023012136 (ebook) | ISBN 9780262547109 (paperback) | ISBN 9780262377294 (epub) | ISBN 9780262377287 (pdf)

Subjects: LCSH: Consciousness.

Classification: LCC BF311 .M4725 2024 (print) | LCC BF311 (ebook) | DDC 153—dc23/eng/20230830

LC record available at <https://lcn.loc.gov/2023012135>

LC ebook record available at <https://lcn.loc.gov/2023012136>